

WHAT IS CLAIMED IS:

1 1. A method for operating a distributed arbiter for a plurality of resources,
 2 comprising:
 3 receiving at one of the resources, requests for the one resource from a plurality
 4 of requesters; and
 5 granting the one resource to one of the requesters according to respective
 6 requester priorities, the respective requester priorities being inversely
 7 related to a number of requests made respectively, by the requesters.

SUB A' 2 2. The method as recited in claim 1 further comprising the requesters
 2 supplying respective requester priority indications to at least the one resource.

1 3. The method as recited in claim 2 wherein requesters respectively
 2 supply number of requests as requester priority indications.

1 4. The method as recited in claim 1 further comprising granting at least
 2 one of the plurality of resources to a requester according to a round robin scheme,
 3 thereby avoiding starvation.

1 5. The method as recited in claim 1 wherein the one resource provides a
 2 grant indication to the one of the plurality of requesters.

1 6. The method as recited in claim 1 wherein the multiple resources
 2 evaluate received requests sequentially.

1 7. The method as recited in claim 6 wherein resources are considered in
 2 an order determined according to resource priority.

1 8. The method as recited in claim 1 wherein the multiple resources
 2 evaluate received requests sequentially, and wherein a requester already granted a
 3 resource removes requests for other resources for arbitration for a next sequential
 4 resource, thereby preventing requesters from receiving multiple grants during an
 5 arbitration cycle.

1 9. The method as recited in claim 1 wherein requester priorities are
2 recalculated after a resource is granted.

1 10. The method as recited in claim 1 further comprising the distributed
2 arbiter performing multiple iterations during one arbitration cycle to allocate
3 resources to requesters.

1 11. The method as recited in claim 9 further comprising requesters
2 recalculating their requests by eliminating requests for a granted resource.

1 12. The method as recited in claim 1 wherein requesters receive multiple
2 grants from multiple resources during an arbitration iteration.

1 13. The method as recited in claim 12 wherein a number of arbitration
2 iterations over multiple arbitration cycles varies dynamically.

1 14. The method as recited in claim 12 wherein the one resource determines
2 which requests to grant in parallel with other resources determining which requests to
3 grant.

1 15. The method as recited in claim 12 wherein the one resource sends a
2 grant indication to a requester being granted its request.

1 16. The method as recited in claim 15 wherein the one resource sends a
2 resource priority indication to the requester being granted its request, indicative of a
3 number of requests received by the resource.

500 A' 1 17. The method as recited in claim 12 wherein a requester selects from
2 among a plurality of grants according to respective resource priorities associated with
3 each granting resource, the respective resource priorities being inversely related to a
4 number of requests made, respectively, for each granting resource.

1 18. The method as recited in claim 17 wherein resource priorities are
2 recalculated each iteration.

1 19. The method as recited in claim 17 wherein requester priorities are
2 recalculated each iteration.

1 20. The method as recited in claim 17 further comprising the requester
2 utilizing a round robin scheme to select from among the plurality of grants.

1 21. The method as recited in claim 20 wherein the round robin scheme is
2 utilized before the requester selects from among the plurality of grants according to
3 the respective resource priorities, the requester selecting from among the plurality of
4 grants according to the respective resource priorities if the round robin scheme does
5 not result in an acceptance of a grant.

1 22. The method as recited in claim 17 further comprising the requester
2 utilizing a random scheme in addition to using calculated priorities to select from
3 among the plurality of grants.

1 23. The method as recited in claim 12 wherein the requester being granted
2 its request sends an accept indication to the resource whose grant it is accepting.

1 24. The method as recited in claim 23 wherein a requester determines
2 which of one or more grants to accept, in parallel with other requesters determining
3 which received grants to accept.

1 25. The method as recited in claim 1 wherein the resources do not transmit
2 grant indications to requesters, each requester determining grant values according to
3 received information from other requesters.

1 26. The method as recited in claim 25 wherein requesters and resources are
2 synchronized in regards to round robin positions.

1 27. The method as recited in claim 2 wherein each of the requesters
2 provides an indication of priority on a bus logically combining the indications of
3 priority, the priority being related to a number of requests being made by each of the
4 requesters.

1 28. The method as recited in claim 27 wherein the indications of priority
2 are unary coded, one bit corresponding to one request.

1 29. The method as recited in claim 28 wherein the bus logically combines
2 the indications of priority.

1 30. The method as recited in claim 28 wherein the indications of priority
2 are a number of requests of respective requesters.

1 31. The method as recited in claim 28 further comprising a requester
2 responding to the priority indication on the bus by not sending a request if the priority
3 indication on the bus indicates a higher priority requester is requesting a resource.

1 32. The method as recited in claim 1 wherein the requesters are nodes of a
2 network coupled to input ports of a network switch and the resources are output ports
3 of the network switch, multiple ones of the output ports being accessible to more than
4 one of the input ports.

1 33. The method as recited in claim 1 wherein the requesters are processors
2 of a multi-processor system and the resources are memories coupled to the processors,
3 each of the memories being accessible to more than one of the processors.

1 34. The method as recited in claim 1 further comprising recalculating
2 requester priorities after each time a resource is granted to a requester.

1 35. An arbitration apparatus for arbitrating requests from a plurality of
2 requesters for a plurality of resources, comprising:
3 means for receiving requests for resources from the requesters; and

means for granting requests according to requester priorities, the requester priorities being inversely related to the number of requests respectively made by the requesters.

36. The arbitration apparatus as recited in claim 35 further comprising:
means for selecting from among a plurality of grants according to resource priorities, the resource priorities being inversely related to the number of requests made for the respective resource.

Sub A' 37. The arbitration apparatus as recited in claim 36 further comprising
means for preventing starvation for grants.

38. The arbitration apparatus as recited in claim 35 further comprising
means for preventing starvation for requests.

39. The arbitration apparatus as recited in claim 35 wherein the arbitration apparatus is one of distributed arbiter and a centralized arbiter.

40. A distributed arbiter comprising:
a plurality of requesters;
a plurality of resources coupled to the requesters through a transport mechanism;
wherein each requester is coupled to provide to each resource requested by the respective requester, a request indication; and
wherein a requested resource is responsive to a plurality of requests to selectively grant one of the requests according to requester priorities, the requester priorities being inversely related to a number of requests being made by respective requesters.

41. The distributed arbiter as recited in claim 40 wherein the requested resource is further responsive to the plurality of requests to selectively grant one of the requests according to a round robin mechanism.

1 42. The method as recited in claim 41 wherein the round robin mechanism
2 is utilized before the requested resource selects from among the plurality of requests
3 according to the requester priorities, the requested resource granting one of the
4 requests according to the requester priorities if the round robin mechanism does not
5 result in a grant.

1 43. The distributed arbiter as recited in claim 40 wherein each requester
2 requesting the requested resource respectively provides a requester priority indication
3 to the requested resource.

1 44. The distributed arbiter as recited in claim 43 further comprising a bus
2 coupled to the requesters and resources, and wherein each of the requesters provides
3 an indication of priority on a bus, the bus logically combining the indications of
4 priority, the priority being inversely related to a number of requests being made by
5 each of the requesters.

1 45. The distributed arbiter as recited in claim 44 wherein the indications of
2 priority are unary coded.

1 46. The distributed arbiter as recited in claim 44 wherein a bus implements
2 a wired-NOR of the indications of priority.

1 47. The distributed arbiter as recited in claim 44 wherein each requester
2 places a unary encoded number of requests on the bus as the indication of priority.

1 48. The method as recited in claim 44 wherein a requester with a lower
2 priority than indicated on the bus is responsive to the priority indicated on the bus to
3 not send its request to a resource.

1 49. The distributed arbiter as recited in claim 40 wherein the requested
2 resource provides a grant indication to a requester whose request is being granted by
3 the resource.

1 50. The distributed arbiter as recited in claim 40 wherein the requested
2 resource is further coupled to provide an indication of resource priority, the resource
3 priority being inversely related to a number of resource requests received by the
4 requested resource.

SUB A' } 1 51. The distributed arbiter as recited in claim 40 wherein a requester is
2 responsive to receiving a plurality of grants from a plurality of resources to accept one
3 of the grants according to priorities of the grants, the priorities being inversely related
4 to a number of requests received by resources.

1 52. The distributed arbiter as recited in claim 51 wherein the requester
2 accepts a grant indicating a particular resource is being allocated to the requester,
3 according to a starvation avoidance mechanism in addition to the priorities.

1 53. The distributed arbiter as recited in claim 40 wherein a requester is
2 responsive to receiving a plurality of grants from a plurality of resources to accept one
3 of the grants according to a random priority determined for the grants.

1 54. The distributed arbiter as recited in claim 40 wherein the requesters are
2 nodes of a network coupled to input ports of a network switch and the resources are
3 output ports of the network switch, multiple ones of the output ports being accessible
4 to more than one of the input ports.

1 55. The distributed arbiter as recited in claim 40 wherein the requesters are
2 processors of a multi-processor system and the resources are memories coupled to the
3 processors, each of the memories being accessible to more than one of the processors.

1 56. A computer program product encoded in at least one computer
2 readable medium to implement an arbitration mechanism to allocate a resource to
3 requesters, the computer program product comprising:
4 a first instruction sequence executable to receive requests for one of the
5 resources from multiple requesters; and

6 a second instruction sequence executable to allocate the one resource to one of
7 the requesters according to requester priorities, requester priorities
8 being inversely related to a number of requests made by each requester
9 for the resources; and

10 wherein the at least one computer readable medium is selected from the set of
11 a disk, tape or other magnetic, optical, or electronic storage medium
12 and a network, wireline, wireless or other communications medium.

1 57. The computer program product as recited in claim 56 wherein the
2 computer program product further includes a starvation avoidance instruction
3 sequence executable to implement a starvation avoidance mechanism.

1 58. A computer program product encoded in at least one computer
2 readable medium to implement an arbitration mechanism to determine which of a
3 plurality of grants received by a requester from a plurality of resources, to accept, the
4 computer program product comprising:
5 a first instruction sequence executable to receive an indication of the grants for
6 respective resources; and
7 a second instruction sequence executable to accept one of the grants according
8 to grant priorities, grant priorities being inversely related to a number
9 of requests received by a resource.

1 59. The computer program product as recited in claim 58 wherein the
2 computer program product further includes a starvation avoidance instruction
3 sequence executable to implement a starvation avoidance mechanism to accept one of
4 the grants.

Add A'